

METHOD AND APPARATUS FOR CLASSIFYING QUERYING NODES

5 CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of U.S. provisional application No. 60/272,887, entitled "Method and Apparatus for Classifying Querying Nodes" (Attorney Docket No. 41966/JEC/X2/134058) filed March 2, 2001, the contents of which
10 are incorporated herein by reference.

FIELD OF THE INVENTION

This application relates generally to Internet multicasting, and more particularly to a method and apparatus
15 for distinguishing different types of querying nodes transmitting multicast group queries.

BACKGROUND OF THE INVENTION

Internet protocol (IP) multicast routers use an Internet
20 Group Management Protocol (IGMP) to learn the multicast group memberships of neighboring nodes, as set forth in Internet Engineering Task Force Request for Comment 2236 entitled "Internet Group Management Protocol, Version 2," November 1997 (hereinafter referred to as RFC 2236), which is incorporated
25 herein by reference. The protocol exchange specified by IGMP to accomplish such learning involves the neighboring nodes sending to the IP multicast router membership reports indicating the multicast groups in which the neighboring nodes participate. The membership reports may either be sent unsolicited or in
30 response to membership queries received from the IP multicast router.

Recently, there has emerged a new class of nodes interested in learning IP multicast group memberships of neighboring nodes. The new class of nodes hereinafter referred
35

to as IP multicast queriers, send IGMP membership queries but do not participate in IP multicast routing protocols. An IP multicast querier may be, for instance, a local area network (LAN) switch that learns the multicast group membership of LAN-attached IP hosts in order to report such multicast groups by proxy to IP multicast routers and to avoid unnecessary flooding of IP multicast traffic received from such IP multicast routers.

With the emergence of IP multicast queriers, a need arises for a way for nodes to, as best as possible, distinguish between two classes of neighboring nodes from which IGMP membership queries may be received, that is, between IP multicast routers on the one hand and IP multicast queriers on the other. Otherwise, neighboring IP multicast routers may be mistaken for IP multicast queriers, or vice versa, resulting in improper withholding or transmission of messages to such neighboring nodes. For example, whereas IP multicast routers are generally required to receive all IP multicast data packets, IP multicast queriers may have no such general entitlement.

SUMMARY OF THE INVENTION

The present invention is directed to classifying querying nodes. The classification information helps distinguish a multicast querier that may receive multicast group membership information but not multicast routing protocol data packets, from a multicast router that may receive both multicast group membership information and multicast routing protocol data packets.

In one embodiment, the invention is directed to a method for classifying a querying node where the method includes receiving a query message from the node, reviewing the query message, and classifying the node as either a multicast router or a multicast querier. The querying node is classified as a

multicast router if the message is a multicast routing protocol query message and the node is either unclassified or classified as a multicast querier. On the other hand, the node is classified as a multicast querier if the message is a multicast group query message and the node is unclassified.

In another embodiment, the invention is directed to a data communication network that includes a plurality of classified nodes and a classifying node with a plurality of ports and one or more databases. The databases include entries associating the classified nodes with respective ports on which messages from the classified nodes were received by the classifying node. The classified nodes include at least one node classified as a multicast querier.

In one embodiment of the invention, the classifying node transmits multicast group membership report messages to the classified nodes via ports associated with the classified nodes.

In another embodiment of the invention, the classifying node transmits multicast routing protocol data packets to a node classified as a multicast router via a port associated with the classified node.

In a further embodiment, the invention is directed to an internet protocol (IP) multicast router that includes a port receiving a query message from a node, a memory storing classification information for the node, and a classification engine coupled to the port and the memory. The classification engine reviews the query message and classifies the node as a multicast router if the message is a multicast routing protocol query message and the node is unclassified or classified as a multicast querier. If the message is a multicast group query message and the node is unclassified, the classification engine classifies the node as a multicast querier.

[illegible]

10

15

20

25

30

35

DETAILED DESCRIPTION OF THE SPECIFIC EMBODIMENTS

FIG. 1 is a schematic block diagram of a data communication
5 network including IP multicast routers 10, 20 and an IP multicast
querier 30 (collectively referred to as IP multicast nodes)
interconnected to each other over cables or other transmission
media, such as wireless, according to one embodiment of the
invention. The multicast routers 10, 20 and querier 30 are each
10 respectively associated with end systems 60, 70, 80, 90, which
are also interconnected to routers 10, 20 and querier 30 over
cables or other transmission media, such as wireless. Routers
10, 20, querier 30 and end systems 60, 70, 80 90 preferably
communicate using a LAN communication medium, such as Ethernet
15 or Token Ring. Ethernet LAN communication media are not limited
to 10 megabit Ethernet, but include other variants, such as Fast
Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet and 802.11b
wireless Ethernet.

The multicast routers 10, 20 may be gateway devices having
20 network interfaces for forwarding multicast communications
originated by the end systems 60, 70, 80, 90 to specified IP
multicast group members. The routers 10, 20 further include
logic for maintaining and exchanging multicast group membership
information with neighboring IP multicast nodes according to the
25 IGMP set forth in RFC 2236.

The multicast querier 30 may be a gateway device such as,
for example, a Layer 2 LAN switch requesting and receiving IGMP
multicast group membership information, but that does not
participate in IP multicast routing protocols. Multicast
30 queriers, however, are not limited to switches, and may include
any device that controls the flow of multicast packets into a
local area network.

35

The end systems 60, 70, 80, 90 may be network end-stations such as, for example, personal computers, workstations, or servers.

FIG. 2 illustrates two nodes in the system of FIG. 1 where one node is a classifying node 100 and another node is a classified node 102. According to one embodiment of the invention, the classifying node 100 is an IP multicast router such as router 10 or 20. The classified node is a node, such as the IP multicast routers 10, 20 or querier 30, transmitting IGMP multicast group query messages 104 to the classifying node 100.

The classifying node 100 classifies the classified node 102 as either an IP multicast querier or an IP multicast router. The classifying node 100 transmits IGMP multicast group reports 106 to the classified node 102 either unsolicited or upon receipt of a multicast group query message. The classifying node 100 also transmits IP multicast routing protocol data packets to the classified node 102 if the classified node 102 is classified as an IP multicast router. IP multicast queriers do not participate in IP multicast routing protocols and do not receive IP multicast routing protocol data packets.

FIG. 3 is a schematic block diagram of an IP multicast router 200 acting as a classifying node according to one embodiment of the invention. The router 200 preferably includes a packet buffer 202, a classification engine 204, one or more databases 206, and an application engine 208. The packet buffer 202 is preferably a FIFO buffer. The classification and application engines 204, 208 may be implemented in software, firmware, and/or hardware. The database 206 may be implemented in a memory 206, such as, for example, a random access memory or the like.

It is understood, of course, that FIG. 3 illustrates a block diagram of an IP multicast router without obfuscating

inventive aspects of the present invention with additional elements and/or components which may be required for the router. These additional elements and/or components, which are not shown in FIG. 3, are well known to those skilled in the art. For example, the router 200 may include a policing engine and a separate packet classification engine in addition to the illustrated components.

10 The router 200 preferably receives inbound packets 210 transmitted by a querying node via its ports (not shown) and stores them in the packet buffer 202. The packets may include IGMP multicast group membership query messages or multicast routing protocol query messages such as, for example, distance vector multicast routing protocol (DVMRP) or protocol independent multicast (PIM) control messages.

15 The packet buffer 202 may include a packet FIFO for receiving and temporarily storing the packets. The packet buffer 202 preferably provides the stored packets or portions thereof to the classification engine 204 for classifying according to conventional mechanisms and/or determining the classification of the querying nodes transmitting the inbound packets 210.

20 The packet buffer 202 may include either or both a header data extractor and a header data cache. The header data extractor may be used to extract one or more fields from the packets, and store the extracted fields in the header data cache as extracted header data. The extracted header data may include, but is not limited to, some or all of the packet header data. For example, the extracted header data may include a query type, an IP multicast group address, and/or the like.

25 The extracted header data is provided in an output signal 216 to the classification engine 204 for processing. The extracted header data may also be provided to the application engine 208 in an output signal 218. In other embodiments, the

output signals 216, 218 may include the whole inbound packet instead of or in addition to the extracted header data.

5 The classification engine 204 classifies the querying node transmitting the inbound packet 210 based on the type of querying message encompassed in the packet. The classification engine further identifies a port on which the inbound packet 210 is received. The port and classification information for the
10 querying node is then stored in the database 206.

The application engine 208 provides application data 214, which may include IGMP multicast group membership report messages and IP multicast routing protocol data, to the packet buffer 202 for transmitting as part of outbound packets 212. According to
15 one embodiment of the invention, IGMP multicast group membership reports are transmitted to the querying node via a port in which its inbound packet was received. The reports are transmitted regardless of the classification of the querying node.

IP multicast routing protocol data packets are transmitted
20 to the querying node if the querying node is classified as an IP multicast router. The packets are transmitted via a port in which its inbound packet was received.

FIG. 4 is a state transition diagram for a querying node according to one embodiment of the invention. The querying node
25 initially begins at an unclassified state 300 where the node is neither classified as an IP multicast querier nor as an IP multicast router. Upon receipt by the classifying node of an initial IGMP multicast group query message 304 from the querying node, the node transitions from the unclassified state 300 to a
30 querier state 302. A querier timer for the querying node is further initialized upon receipt of the query message. Receipts of additional multicast group query messages 306 from the querying node retains it in the querier state 302 and causes resets of the querier timer. If the querier timer times-out 308

35

without the classifying node receiving a multicast group query message from the querying node, the querying node is declassified and the node transitions to the unclassified state 300.

If the classifying node receives a multicast routing protocol query message (router query) 310, 312 from the querying node that is either unclassified 300 or classified in the querier state 302, the node transitions to a multicast router state 314.

A router timer for the querying node is also initialized upon receipt of the router query. Such router queries include but are not limited to DVMRP or PIM control message packets.

The querying node, while classified in the multicast router state 314, receives IGMP membership reports as well as other IP multicast data packets from the classifying node. Receipts of additional router queries 316 or multicast group queries 318 from the querying node retains the querying node in the multicast router state 314. According to one embodiment of the invention, receipt of router queries 316 from the querying node causes resets of the associated router timer. Receipt of multicast group queries 318 while classified in the multicast router state 314, however, does not reset the router timer or the querier timer. If the router timer times-out 320 without the classifying node receiving a router query from the querying node, the querying node transitions to the querier state 302.

FIG. 5 is a flow diagram of a process for classifying querying nodes according to one embodiment of the invention. The process starts, and in step 400, the classifying node determines if a query message packet has been received. If the answer is YES, a determination is made as to the type of the packet. If the packet is a multicast group query packet, as determined in step 402, and the querying node is either unclassified or already classified as a multicast querier node, as determined in step 404, the node is classified or reclassified as a multicast

querier node in step 406. The querier timer associated with the querying node is also initialized or reset in step 408. In step 5 410, the port in which the packet was received is identified. In step 412 the database 206 is updated to reflect that the node transmitting via the identified port is a querier node. IGMP multicast membership group report messages may then be transmitted to the querying node via the identified port.

10 If, however, the querying node has been classified as a multicast router node upon receipt of the multicast group query packet as determined in step 416, the node maintains its current classification state. Thus, the node continues to be classified as a multicast router node in step 407 without resets in either 15 the router timer or the querier timer. The port receiving the multicast group query packet is further identified in step 410 and the database updated accordingly in step 412.

If instead of a multicast group query packet the classifying node receives a multicast routing protocol query 20 message packet as determined in step 418, the router timer is initialized or reset in step 426. In step 407, the querying node is classified or reclassified as a multicast router node. This occurs whether the node is currently unclassified, classified as a multicast querier node, or classified as a multicast router 25 node. In step 410, the port receiving the routing protocol query message packet is identified, and database updated accordingly based on the information in step 412. Future multicast routing protocol data packets may then be transmitted to the querying node via the identified port.

30 FIG. 6 is a timer processing flow diagram for the querier or router timer according to one embodiment of the invention. The process starts, and in step 430, the classification engine 204 determines whether a timer being monitored has timed out. If the answer is YES, and the querier timer has timed-out as

35

determined in step 432, the node is declassified in step 434. Accordingly, database 206 is also updated by in step 440 by
5 deleting the classification and port information for the node.

If the router timer has timed-out as determined in step 436, the node is reclassified as a multicast querier in step 438 and the querier timer reset in step 442. The database entry for the node is also updated with the new classification information
10 in step 440.

FIG. 7 is a layout diagram of the database 206 of FIG. 3 maintained by the classifying node according to one embodiment of the invention. The database preferably includes entries associating the querying nodes with ports on which query messages
15 from these nodes were received. According to the illustrated embodiment, each entry includes a source address 500, port number 502, classification 504, and virtual local area network (VLAN) number 506. The source address 500 corresponds to the source address of the querying node. The port number 502 preferably
20 identifies a port on the classifying node via which a query message transmitted by the querying node was received. Such query messages may include multicast group queries or multicast router queries. The VLAN number 506 identifies a VLAN on which the query message was received. The classification 504 field
25 identifies a classification of the querying node as determined by the classification engine 204.

According to one embodiment of the invention, multicast membership reports are transmitted to a querying node via its associated port identified in the database 206. IP multicast
30 routing protocol data packets, however, are transmitted to ports associated with nodes that have been classified as a multicast router or have otherwise, through IGMP, requested a particular data stream.

Although this invention has been described in certain
specific embodiments, those skilled in the art will have no
5 difficulty devising variations which in no way depart from the
scope and spirit of the present invention. It is therefore to
be understood that this invention may be practiced otherwise than
is specifically described. Thus, the present embodiments of the
invention should be considered in all respects as illustrative
10 and not restrictive, the scope of the invention to be indicated
by the appended claims and their equivalents rather than the
foregoing description.

15

20

25

30

35